Demand for Self Control: A Model of Consumer Response to Programs and Products that Moderate Consumption

Nathan Berg* and Jeong-Yoo Kim

Abstract: Is it better to apply effort to increase personal consumption, or control what one wants? The model presented here provides a characterization of demand for self control, namely, its responsiveness to price and risk. Unlike most other models of self control, the model does not identify self control with time inconsistency or rely on the multiple-selves framework. Self control refers to resources allocated to preference transformation technology enabling consumers to moderate desire for ordinary consumption by reducing threshold levels required to achieve goals or target-levels of consumption. Consumers face a choice between allocating resources toward increasing expected levels of consumption or increasing chances of contentment through self control. Because of strong income effects, demand for self control turns out to be non-monotonic in price and sometimes discontinuous, revealing potential for unanticipated and sometimes surprisingly large responses to small changes in price. The model is used to analyze consumers' willingness to follow new regulations, take up credit counseling, enroll in financial literacy programs, and purchase products aimed at improving financial decision making through cultivation of self control.

Keywords: Preference Choice, Preference Change, Moderation, Restraint, Desire, Financial, Decision Making, Consumer Credit, Consumer Finance, Institutional Design, Ecological Rationality, Bounded Rationality, Behavioral Economics

JEL Codes: D03, D60, D11, D13, D18, G21, G28, C65, B30, B52

---

* *Berg:* Associate Professor of Economics, School of Economic, Political, and Policy Sciences (EPPS), University of Texas at Dallas, GR 31 211300, Box 830688, Richardson, Texas 75083-0688. Email nberg@utdallas.edu.

Kim : Professor of Economics, Kyung Hee University, Phone and Fax +822-961-0986. Email jyookim@khu.ac.kr.

The authors thank Michael Haydel, Fatma Sardina, Zhengzheng Wang, and Nohel Zaman for valuable research assistance.

1 Introduction

"Excessive" borrowing, housing demand and risk taking are frequently mentioned by commentators attempting to explain recent crises in US financial and housing markets. The notion of excessive consumption (especially of so-called "toxic" loans and financial products) is the maintained premise that motivates various arguments, put forward from distinct political viewpoints, for regulatory reforms, new programs, and products aimed at improving consumers' self control.[1] Different interpretations of recent events (in financial, housing—and now labor—markets) put varying degrees of emphasis on the role that flawed consumer decision making has played, alternatively focusing on consumers' excessive borrowing (Holzer, 2009), regulatory failure (Brown, 2007; Bernanke, 2010), insufficient fiscal stimulus (Krugman, 2009; 2010), and other institutional failures that may have encouraged accounting fraud (Black, 2005; Akerlof and Romer, 1993) as root causes. Rather than address controversies surrounding these interpretations, this paper takes for granted that consumers are capable of excess, assuming that they reflect on this very possibility, and then attempt to change preferences in a purposeful and effortful manner using technology that aids self control.

This paper introduces a simple model of self control that, unlike most of the self control literature in economics, does not rely on time inconsistency or the multiple-selves framework. The model introduced here is used to analyze how consumers

[1] Excessive consumption, or "over-consumption," and its welfare consequences is discussed in reports by the Financial Crisis Inquiry Commission (2010), O'Donoghue and Rabin (2002), Bae and Ott (2008), and numerous others in behavioral economics. See, however, Sugden (2004), Starr (2007), and Berg and Gigerenzer (2010) for different views on the normative authority of standard neoclassical rationality axioms and the methodological implications of modelers who interpret agents in their models as engaging in excessive consumption. The model presented here describes an agent who, after thinking through different ways of achieving contentment under different future contingencies, regards some levels of acquisitiveness as excessive and uses preference change technology to moderate his or her desire for ordinary consumption.

might respond to new regulatory institutions, educational campaigns, and financial services that aim to help consumers improve self control and moderate potentially excessive demand for certain goods and services.[2] The model draws motivation from the observation that, after introspecting on one's goals and various means of reaching them, consumers routinely wish to modify those goals in a way that makes them more achievable (i.e., less prone to the discontentment of not having reached the minimum level of ordinary consumption used by the consumer to benchmark his or her goals). This motive leads consumers to substitute out of ordinary consumption and into a distinct class of expenditures on preference-change technology (i.e., self control), which reduces desire for ordinary consumption. Acquisition of costly preference-change technology is referred to as demand for *self control,* in distinction with expenditures on everything other than self control, which is referred to as *ordinary consumption* (i.e.,

[2]The model abstracts from institutional detail specifying particular forms of self control technologies. Aggregated as a single category of resource expenditure, the consumer's allocation of resources into self control can be interpreted flexibly, for example, as taking up financial literacy programs, credit counseling, or choosing privately marketed products that offer intentionally restricted choice sets that serve as commitment devices. Such behavior finds motivation in the finding of widespread financial illiteracy (Lusardi and Mitchell, 2006; Huston, 2010), and expanding such programs is explicitly recommended by Yao, Hanna and Lindamood (2004), Emmons (2005) and many others. Products and programs that make self control technology more available are generally interpretable in our model as a reduction in the price of self control. Examples include the Save More Tomorrow program (Thaler and Benartzi, 2004); greater availability of information making it easier for consumers to be savvy price shoppers on the lookout for favorable borrowing terms (Kerr and Dunn, 2008); consumers' use of debit cards as a device to impose spending limits on themselves (Stavins, 2008); tax withholding as a savings commitment device (Barr and Dokko, 2008); borrowing limits on credit cards (Dey and Mumy, 2005); social security as a forced savings device that aids self control (Kumru and Thanopoulos, 2008); and smaller package sizes that might, for example, help smokers encourage themselves to keep open the option of quitting by buying packs instead of cartons (Wertenbroch, 1998), although small package sizes sometimes appear to have perverse effects that increase consumption (do Valle, Pieters, and Zeelenberg, 2008). Even credit cards can, in some contexts, be considered as a self control technology, providing a means to counter the temptation to sell illiquid retirement assets in tax deferred savings accounts which would incur large penalties (Bertaut, Haliassos and Reiter, 2009). Important differences in financial literacy and other consumer behaviors have been documented in low-income populations (Servon and Kaestner, 2008; Beverly and Sherraden, 1999), which translates in our model to distinct wealth endowment and consumption-aspiration parameters.

goods and services that monotonically increase utility, holding preferences fixed).

The exogenous preference model of neoclassical economics taught in most microeconomics textbooks focuses on maximization of satisfaction (i.e., a utility function that depends on ordinary consumption) while holding preferences and the exogenously given resource constraint fixed. An alternative goal-seeking model that fits squarely within the constrained optimization methodology of neoclassical economics (yet, puzzlingly, is rarely, if ever, addressed in this literature) can be described as follows. Given the flows of ordinary consumption that one currently owns, maximize the experienced satisfaction derived from it, by engaging in purposeful preference change with the aid of preference-change technology.[3] The consumer asks him or herself: Is it better to increase expected ordinary consumption conditional on what one wants, or control what one wants conditional on what one has?[4]

Following standard competitive price theory, the model in this paper assumes the consumer takes the price of preference-change technology as given (relative to ordinary consumption). With standard assumptions, the model produces an optimal choice rule for a consumer who simultaneously considers benefits and costs of allocating

[3] Goal seeking is modeled and investigated empirically by Güth (2010), Bergninghaus, Güth, Levati and Qui (forthcoming), and Güth , Levati and Ploner (2009; 2010). Fisher and Montalto (2010) use a reference-point-dependent model similar in spirit to ours, combined with empirical analysis of a large survey data set, to argue for the empirical relevance of reference points, goal setting, and the importance for policy makers to consider multiple consumer goals. Jain (2009) presents a related model of goal seeking but without the contentment threshold built into the utility function in our model. Consistent with this paper, Samwick (2006) documents heterogeneous goals among savers; Zhang, Huang and Broniarczyk (2010) show additional dimensions of the importance of consumer goals; and a number of experimental studies lend support to goal setting as a primary behavioral mechanism resulting in contentment and discontentment (Lee and Ariely, 2006; Trudel and Cotte, 2008; Drolet, Luce and Simonson, 2009).

[4] Starkly different views about how humans achieve happiness can be found in various philosophical traditions, discussed in detail in Bruni and Porta (2007), who emphasize social ties as sources of satisfaction and the need to learn special techniques for moderating one's desires. Although the topics of self control and moderation of desire are sometimes described in esoteric terms, this paper focuses on self control in the context of the ongoing financial market crisis and numerous proposals being debated to help consumers manage their borrowing, consumption of housing, and use of financial services.

one's resources (effort or wealth endowment) to the expected quantity of ordinary consumption versus preference-change technology that increases the likelihood that one is satisfied with the quantity of ordinary consumption he or she will receive ex post.[5] With regard to the role of consumer behavior and consumer finance products in recent housing and financial crises, ordinary consumption should be interpreted as a composite of housing, credit, and financial products whose status as "goods" has come under question (e.g., payday lending services).

Policies, programs and products aimed at cultivating self control are interpreted in this model as interventions that lower the price of self control. The model then provides an answer to the question of how consumers will respond to new incentives to take up self control. The model is simple. Self control directly raises utility by transforming preferences such that every fixed level of ordinary consumption yields greater experienced satisfaction because, ex ante, it carries a reduced risk of falling short of one's goal. This mechanism provides a simple motive for self control and facilitates applied analysis relevant to consumer finance and the uptake of services for cultivating self control, without relying on time inconsistency or the multiple-selves framework (studied with great insight, and both theoretical and empirical rigor, by Tirole, 2002; Gifford (2002); Loewenstein, Read and Baumeister, 2003; Bénabou and Tirole, forthcoming; and others who focus on time inconsistency as a means of modeling problems in self control).

[5]Different interpretations of the consumer's resource allocation problem enjoy empirical support, alternatively, in units of effort or money. For example, Oaten and Cheng (2007) show that simple practice (i.e., effort) can have large effects on self control in the financial decision making context. And if there were any doubt that demand for self control responds systematically to its price in units of money, a rather impressive empirical record establishes large demand-side effects of price effects on self control (Feldstein, 1995; Ameriks, Caplin, Leahy and Tyler, 2007; Duygan-Bump and Grant, 2008), although it remains unclear whether the measured magnitudes can be extrapolated to new contexts. Self control in the context of developing economies is studied by Semboja (2004) and Ashraf, Karlan and Yin (2005).

Even in the simplest version of our model, several surprising results emerge. First, self control is an inferior good, with demand decreasing as a function of an exogenously given wealth endowment. Second, because of the joint interaction of ordinary consumption and the contentment threshold (a choice variable under the decision maker's control), demand for self control exhibits unusually strong income effects, resulting in upward-sloping demand curves on a dense subset of the model's parameter space. Thus, the model provides a new explanation (other than the commitment device hypothesis found in multiple-selves models) for consumers who prefer high-priced self-control technologies. Finally, because demand for self control is non-monotonic in price, and because the globally maximizing level of self control switches discontinuously from interior to corner solutions, the model's price and demand dynamics are highly unstable. Small changes in price (over large regions of the model's parameter space) produce dramatic shifts in demand for self control, perhaps equivalent to differences in social norms about personal savings in the U.S. and Japan, or religious revolutions that focus on moderation of desire for material or worldly consumption. Therefore, attempts by policy makers, nonprofits, banks, and other financial services firms to reach consumers using standard economic incentives (i.e., reductions in the price of self control) may have no effect at all, moderate effects, huge effects, or (in the case where the decision making environment is situated on an upward-sloping portion of the demand schedule for self control) large effects in the opposite direction as had been anticipated. Unstable demand for self control (again, due to non-monotonicity and occasional discontinuities) is, in essence, an anything goes result. Given a detailed description of exogenous factors in the environment, the model does indeed provide (mostly) unambiguous predictions about consumer responses to incentives aimed at inducing greater self control. However, any imprecision in estimates describing the decision making environment can lead to unanticipated demand responses.

One novel element of the model is that self control is a technology that directly reduces reference-point levels of ordinary consumption needed to achieve psychological contentment (e.g., everything from credit counseling and financial literacy programs to religious teaching or philosophical practice aimed at moderating demand for ordinary consumption). This implies that preferences are, to a partial extent, a choice variable (c.f., George, 2001).

In this sense, ours can be described as a contentment-aspiration theory of self control, modeled after extensive theoretical work and empirical evidence in favor of satisficing (Güth, 2010; Güth, Levati and Ploner, 2009; 2010; Bergninghaus, Güth, Levati and Qui, forthcoming). Gifford (2002) presents an original dynamical model where self control is modeled as inhibition of emotion, which is needed to undertake deliberation about the decision making environment. Most theoretical models of self control in behavioral economics identify self control problems with dynamic inconsistency (e.g., Thaler and Shefrin, 1981; Carrillo, 1998; Kim, 2005; Fudenberg and Levine, 2006). The one-period model introduced in this paper complements this line of work by showing a distinct motive for self control, interpreted as purposeful moderation of one's preferences through a mechanism of contentment aspiration or goal seeking.

Research on self control in a consumer finance context includes credit cards (Castronova and Hagstrom, 2004; Kawaja, 1969), retirement saving (Thaler and Shefrin, 1981; Laibson, Repetto and Tobacman, 1998; Otto, Davies and Chater, 2007; Benartzi and Thaler, 2007), other investment decisions (Hilton, 2001; Bergstresser and Porterba, 2004; Lai, 2006; Chatterjee, Goetz and Palmer, 2009), personal bankruptcy (White, 2009), smoking (Kan, 2007; Ernst, Hogan, Vallas, Cook and Fuller, 2009), food choice (Smith, 2002; Lin and Yen, 2008), and other consumer decisions (Sally, 2000; Karlsson, 2003; Ramanathan and Williams, 2007; Vohs and Faber, 2007; Fryer, 2010; Laran, 2010; Hastings and Washington, 2010). Self control, restraint and moderation technologies are studied in economic history as well (Gamber, 2003; Zhang, Huang and Broniarczyk, 2010), which motivates some interpretations of our model mentioned above.

The paper proceeds as follows. Section 2 specifies the decision problem and solves for optimizing quantities of ordinary consumption and self control, which are expressed explicitly as functions of the price of self control, parameters specifying self control technology, production technology, and the wealth or resource constraint. Section 3 reports comparative statics showing that self control technology in the standard consumer problem is an inferior good. Section 4 depicts demand curves for self control as a function of price, revealing strong income effects that produce upward sloping, and sometimes discontinuous, self control. Section 5 concludes with a discussion of how this extremely simple model illuminates the complexities of consumer finance and institutional design.

2 The Model

Let X represent the uncertain quantity of ordinary consumption the decision maker receives ex post, with cdf $F(x)$, density $f(x)$, and expected value $\mathrm{E}[X] = \mu$. The function $\mu(a) = \mathrm{E}[X]$ describes how the choice variable a (for acquisitiveness), representing allocation of effort into acquiring ordinary consumption, affects the consumer's expected payoff in units of ordinary consumption. We assume $\mu(a)$ is increasing and differentiable (i.e., $\mu'(a) > 0$).[6]

The other effortful activity in the model aside from ordinary consumption that

[6]The following paragraph departs from standard consumer theory by modeling a consumer's explicit reasoning about allocating scarce resources to a preference-transforming activity which changes the expected utility associated with any fixed quantity of ordinary consumption.

increases utility is denoted s (for self control). Effort toward self control, s, affects utility by reducing the probability of falling short of one's goal or target level of ordinary consumption, all else equal.[7] All else is not equal, however, when the decision maker adjusts s. Raising s (to a higher level of effort toward self control) comes at the real cost of reduced a and, consequently, a reduced level of expected ordinary consumption, $\mu(a)$. At the cost of effort that would otherwise go toward ordinary consumption, the decision maker in our model can choose instead to spend some or all of that effort lowering the threshold that defines contentment and thereby actively transform his own preferences.[8] The mechanism by which s moderates desire is assumed to take the specific form of a controllable threshold-level of ordinary consumption. Contentment and discontentment are psychological states defined in terms of this threshold parameter, based on Simon (1956), Gigerenzer and Selten (2002), and Güth's (2010) notion of satisficing.[9]

The threshold that defines contentment is denoted $t = t(s, t_0)$, which depends on self control s and an exogenously given initial threshold t_0. If $X < t(s, t_0)$, then

[7] Preference-transformation technology, or "self control" for short, enables one to feel content with less. Discontentment and contentment refer to a psychological component of the utility function defined explicitly below. Allocation of effort for the purpose of moderating demand for ordinary consumption falls under the heading of self control, which includes formal market transactions for books, services, and experiences that can be regarded as self control technology and therefore not considered ordinary consumption. Ordinary consumption, then, is a residual consumption category, including all goods and services not used for self control.

[8] Preference transformation technologies put forward in the self control literature include Wertenbroch (1998); Thaler and Benartzi (2004); Dey and Mumy (2005); Kumru and Thanopoulos (2008); do Valle, Pieters, and Zeelenberg (2008); Barr and Dokko (2008); Kerr and Dunn (2008); Stavins (2008); and Bertaut, Haliassos and Reiter (2009). An alternative parameterization of self control will be introduced in a subsequent section. In the simplified linear version of the general model developed there, it is most natural to apply the label "self control" not to s but to the resulting distance by which preferences are transformed (in a sense that will be made concrete shortly).

[9] Jain (2009) studies goal attainment, with the interesting finding that it is sometimes better for consumers not to have any goals. Our model is not a satisficing model because it wholly adopts the assumption of expected utility maximization and its implication of exhaustive search through the choice set rather than limited search as in the satisficing model, although the threshold mechanism that defines discontentment functionally resembles satisficing.

the event of discontentment occurs and the psychic cost of discontentment reduces utility. When $X \geq t(s, t_0)$, the discontentment term is zero. For low values of t, discontentment can be interpreted as personal bankruptcy. Very large values of t indicate intense desire for ordinary consumption in the form of a large requirement to avoid the psychic cost of having not achieved one's goal.[10]

Because s reduces the threshold separating the range of ordinary consumption levels into discontented and contented outcomes, a natural assumption about the preference-transformation mechanism is

$$\frac{\partial t(s, t_0)}{\partial s} < 0. \tag{1}$$

Effort allocated to self control reduces the level of ordinary consumption needed to avoid the psychological state of discontentedness. It is similarly natural to normalize units of self control such that $t(0, t_0) = t_0$, implying that the default threshold t_0 is the relevant reference-point level of consumption whenever zero effort is allocated to self control.

The magnitude by which ordinary consumption falls short of the threshold t is $\max\{t - X, 0\}$ and utility is then specified as the function $u(X, \max\{t - X, 0\})$, which is increasing in the first argument and decreasing in the second. The key trade-off in the model concerns the reduction in effort applied to acquisitiveness (a) and consequently lower expected ordinary consumption X that results, whenever an additional unit of effort is applied toward s to lower contentment threshold t.

[10] DeJong and Ripoll's (2006) alternative approach to specifying the utility function to model self control adds a temptation cost to the objective function that is proportional at each period to the utility that would have been derived from spending one's entire lifetime wealth on present consumption.

2.1 From Units of Effort to Units of Ordinary Consumption

In this section, we translate the consumer's effort variables a and s to units of ordinary consumption, which allow us to quantify the opportunity cost of self control in units of expected ordinary consumption. First, we identify self control with the distance by which the contentment threshold is reduced from t_0. This introduces an alternative definition of self control in units of ordinary consumption:

$$\delta = \delta(s) \equiv t_0 - t. \tag{2}$$

According to this definition, the non-negative δ measures the decrease from t_0 that produces the moderated, or self controlled, threshold $t = t_0 - \delta$. The idea that effort applied to self control monotonically lowers the contentment threshold is formalized by the assumption:

$$\frac{\partial \delta(s)}{\partial s} > 0. \tag{3}$$

The normalization $t(0, t_0) = t_0$ is equivalent to

$$\delta(0) = 0. \tag{4}$$

The goal of this section is to compute $\frac{\partial \mu(a)}{\partial \delta}$ along the resource constraint so that its magnitude can be parameterized as the price of self control (i.e., the ordinary consumption foregone for one unit of self control). We take this in three steps.

First, to measure the additional s required to increase δ by one unit (i.e., reduce t by one unit) given by preference-change technology $\delta(s)$, we need the inverse relationship δ^{-1} mapping δ into s, whose first derivative is expressed in differential form as $\frac{\Delta s}{\Delta \delta} = \frac{\partial \delta^{-1}}{\partial \delta}$.

The aggregate effort/resource constraint plays an important role in the model, given by the effort aggregation function $e(a, s)$, which is assumed to be smooth, differentiable and strictly increasing in both arguments, and the parameter A representing the consumer's endowment of aggregate effort/resource (given exogenously by biological, social, institutional, or historical factors) as:

$$e(a, s) \leq A. \tag{5}$$

Along the manifold where this constraint binds, there is an implicit tradeoff between self control and expected ordinary consumption. The second step en route to computing the price of self control is to express the cost of one unit of s in terms of a, which can be computed (along the manifold defined by the set on which the effort constraint holds with equality) as:

$$\frac{\Delta a}{\Delta s} = -\frac{\partial e(a, s)}{\partial s} / \frac{\partial e(a, s)}{\partial a}. \tag{6}$$

Finally, the third step requires a calculation of the opportunity cost (given by the production technology $\mu(a)$) of δ in terms of μ. The sensitivity of μ to a, $\frac{\partial \mu(a)}{\partial a}$, can also be represented as the ratio of differentials $\frac{\Delta \mu}{\Delta a}$.

Combining the three terms introduced above produces an expression for the price of self control in units of foregone expected consumption:

$$\frac{\Delta \mu}{\Delta a} \frac{\Delta a}{\Delta s} \frac{\Delta s}{\Delta \delta} = \frac{\partial \mu(a)}{\partial a} \left(-\frac{\partial e(a, s)}{\partial s} / \frac{\partial e(a, s)}{\partial a}\right) \frac{\partial \delta^{-1}}{\partial \delta}. \tag{7}$$

2.2 The Decision Problem

The decision maker's ex ante problem can now be specified in detail. The decision maker chooses a and s to maximize expected utility subject to three constraints consisting of production technology, self-control technology, and the effort resource constraint:

$$\max_{(a,s) \in \Re^2_+} \mathrm{E}[U(X, \max\{t_0 - \delta(s) - X, 0\})], \text{ such that } E[X] = \mu(a),\ \delta = \delta(s), \text{ and } e(a, s) \leq A. \tag{8}$$

As long as the production and self-control technologies are defined by one-to-one (i.e., invertible) functions so that $a(\mu) = \mu^{-1}$ and $s(\delta) = \delta^{-1}$ exist, the choice variables can be re-specified in terms of μ and δ by inserting these inverse relationships into the resource constraint:

$$\max_{(\mu,\delta)\in\Re^2_+} \mathrm{E}[U(X, \max\{t_0 - \delta - X, 0\})], \text{ such that } e(a(\mu), s(\delta)) \leq A. \tag{9}$$

As the next section shows, with appropriate re-scaling, the effort constraint in (9) can, once again, be re-expressed as a budget constraint in units of ordinary consumption, where the opportunity cost in equation (7) is re-parameterized as p, which represents the price of δ in terms of μ.

2.3 Linear Version of the Model

Suppose the production in ordinary consumption is linear in a:

$$\mu = \mu_0 + v_a a, \Rightarrow \frac{\partial \mu}{\partial a} = v_a. \tag{10}$$

Suppose that the self-control technology is described by a linear relationship between δ (the *change* in the contentment threshold) and s:

$$\delta = v_s s, \Rightarrow \frac{\partial \delta^{-1}}{\partial \delta} = 1/v_s. \tag{11}$$

There is no constant in the linearization above because of the normalization in (4) based on the requirement that, when zero effort is applied to self control, the change in the contentment threshold is also zero. The linear version of the model also supposes that the effort resource constraint is linear:

$$a + s = A, \Rightarrow -\frac{\partial e(a,s)}{\partial s} / \frac{\partial e(a,s)}{\partial a} = -1. \tag{12}$$

It follows from linearization of the three constraints that $\frac{\Delta\mu}{\Delta a}\frac{\Delta a}{\Delta s}\frac{\Delta s}{\Delta \delta} = v_a(-1)(1/v_s)$. Thus, raising δ (i.e., reducing t) by one unit will cost v_a/v_s in units of μ: $\frac{\partial \mu}{\partial \delta} = -v_a/v_s$. We refer to the magnitude of this opportunity cost as the price of self control:

$$p = v_a/v_s. \tag{13}$$

2.4 The Effort Constraint is a Budget Constraint

With linearized functional forms, the effort constraint can be re-written in the easy-to-recognize form of a budget constraint. Multiplying the equation $a + s = A$ from (12) through by v_a and adding μ_0 to both sides obtains (after substituting $\mu = \mu_0 + v_a a$ on the left-hand side):

$$\mu + v_a s = v_a A + \mu_0. \tag{14}$$

The expression on the right-hand side represents the value of the aggregate effort endowment A in units of ordinary consumption, which is equal to the maximum that μ can possibly take on (i.e., allocating all effort to acquisitiveness), denoted $M \equiv v_a A + \mu_0$. Finally substituting $p\delta = (v_a/v_s)(v_s s) = v_a s$ produces the budget constraint:

$$\mu + p\delta = M. \tag{15}$$

Equation (15) translates choice variables a and s in the effort constraint $e(a, s) = A$ to a straightforward budget constraint in the transformed choice variables μ and δ, where parameter p represents the price of self control.

2.5 Additively Separable Utility and Discontentment

Three further simplifications are introduced before solving the decision maker's constrained optimization problem and analyzing the statics of demand for self control as a function of exogenous variables given by the decision maker's environment. The specification below assumes that U is additively separable in ordinary consumption and discontentment. Second, the specification assumes, apart from the discontentment term, that preferences over distributions of X are risk neutral. Third, the specification ignores the magnitude by which ordinary consumption falls below the threshold t and represents discontentment as a binary indicator variable. Binary discontentment could derive from a psychological process by which failing to meet one's goal imposes a fixed psychic cost. A long-lasting malaise after failing to be admitted to any select colleges would be one example. The sting of rejection after failing to attract a marriage partner that matches one's aspiration would be another. Discontentment in this model can be interpreted as a change in mood upon the event of having not reached one's goal that sours all other consumption experiences by a fixed amount. The resulting utility specification is:

$$U(X, \max\{t - X, 0\}) = X - \beta \mathbf{1}(X < t), \tag{16}$$

where $\beta > 0$ is the model's only preference parameter, which measures the extent to which the decision maker places weight on the possibility of falling below the contentment threshold when weighing different choices of δ and μ and the associated distributions of U. Note that this utility specification is a special case of the general specification above (9).

Substituting $\mu = M - p\delta$ from the budget constraint (15) simplifies the constrained optimization problem in two choice variables to an unconstrained optimization problem in δ, with the following objective function:

$$v(\delta) = M - p\delta - \beta F(t_0 - \delta). \tag{17}$$

Under the expectations operator, X is mapped into μ, which is then substituted as a function of δ, and the indicator function that depends on X is mapped into its expectation: $\mathrm{E}[\mathbf{1}(X < t)] = F(t) = F(t_0 - \delta)$.

The objective function $v(\delta)$ explicitly shows the trade-off that is central to the model. Self control as measured by the decision maker's choice of δ provides the benefit of reduced risk of discontentment (i.e., reducing t from its initial value of t_0 to the more moderate or restrained goal, $t_0 - \delta$), but comes at the cost of a reduced level of expected ordinary consumption (i.e., reducing its expected value from μ from

M to $M - p\delta$). When zero effort is allocated to self control ($\delta = 0$) and 100 percent is allocated to acquisitiveness, the objective function simplifies to $v(0) = M - \beta F(t_0)$. The only way that self control can provide improvements in utility is by reducing the β-weighted risk of discontentment more than the reallocation of effort into self control costs.

It would be a mistake to compute the first-order condition for an interior maximizer based on differentiation of (17) with respect to δ using the notation given there because of a subtle dependence of $F(x)$ on δ. As δ changes, so too does the distribution function $F(x)$, even holding the argument x constant. The indirect dependence of F on μ not explicit in the notation in (17) stems from the fact that δ affects μ through the budget constraint ($\mu = M - p\delta$). In the objective function (17) in which the explicit argument of F also depends on δ, two additional terms are needed to differentiate $F(t_0 - \delta)$ with respect to δ:

$$\frac{dF(x)}{d\delta} = F'(x)\frac{\partial x}{\partial \mu} + \frac{\partial F(x)}{\partial \mu}\frac{\partial \mu}{\partial \delta} = f(t_0 - \delta)(-1) + \frac{\partial F(t_0 - \delta)}{\partial \mu}(-p). \tag{18}$$

By specifying $F(x)$ such that μ is strictly a position parameter that shifts the distribution function to the right or left (e.g., if X is normally distributed), the dependence of F on μ can be expressed directly in the argument of F.

We now assume that X is normally distributed with exogenously given $\mathrm{Var}(X) = \sigma^2$, which allows the cdf to be rewritten in terms of the standard normal cdf $\Phi(z)$, and pdf $\phi(z)$, both of which are independent of μ and σ:

$$F(t) = F(t_0 - \delta) = \Phi(\frac{t_0 - \delta - \mu}{\sigma}) = \Phi(\frac{t_0 - (1-p)\delta - M}{\sigma}), \tag{19}$$

where the last equality follows from substitution of $M - p\delta$ for μ from the budget constraint. After substituting away all channels of dependence on μ so that the objective function is solely a function of δ, $v(\delta)$ can be rewritten as:

$$v(\delta) = M - p\delta - \beta\Phi(\frac{t_0 - (1-p)\delta - M}{\sigma}). \tag{20}$$

For simplicity, the exogenous parameters can be stacked in the vector $\theta \equiv [p, t_0, M, \beta, \sigma]$ and the global maximizer of (20) is denoted δ^*, or $\delta^*(\theta)$ to make its dependence on the exogenous parameters explicit. It turns out that corner solutions to the problem of maximizing $v(\delta)$ with respect to δ occur frequently, and expressions will be presented below stating where zero self control, $\delta^* = 0$, and maximum self control, $\delta^* = (M - \mu_0)/p$, are global maximizers. For an interior maximizer δ^*, $0 < \delta^* < (M - \mu_0)/p$, the first-order condition for a maximizer of (20) is:

$$-p + \beta\phi(\frac{t_0 - (1-p)\delta^* - M}{\sigma})\frac{1-p}{\sigma} = 0. \tag{21}$$

The second-order condition for an interior maximizer is:

$$-\beta\phi'(\frac{t_0 - (1-p)\delta^* - M}{\sigma})(\frac{1-p}{\sigma})^2 < 0, \tag{22}$$

which requires the pdf of X to be strictly increasing at $t_0 - \delta^*$ or, equivalently, that the argument of $\phi(\cdot)$ is negative:

$$t_0 - (1-p)\delta^* - M < 0. \tag{23}$$

To solve (21) explicitly for δ^* requires taking the inverse of $y = \phi(z) = \frac{1}{(2\pi)^{1/2}} e^{-\frac{1}{2}z^2}$. The range of y values that this function hits is the half-open, half-closed real line segment $(0, \frac{1}{(2\pi)^{1/2}}]$. The inverse relation can be expressed as $z = \pm[-2\log((2\pi)^{1/2}y)]^{1/2}$. By virtue of the second-order condition, however, only the negative term above is needed, which we use to re-express the first-order condition (21):

$$\frac{p}{1-p}\frac{\sigma}{\beta} = \phi(\frac{t_0 - (1-p)\delta^* - M}{\sigma}). \tag{24}$$

This representation of the first-order condition can be interpreted as choosing self control in a manner that equates the risk-benefit-weighted price of self control (the left-hand side) with its marginal benefit in terms of reduction in the probability of discontentment (on the right-hand side).

The inverse of ϕ is, in general, not a function but a set-valued relation. However, the inverse $\phi^{-1}(\cdot)$ can be functionally identified with the negative element from among the pair of elements in its image, enabling us to apply the inverse function to both sides of (24):

$$-[-2\log((2\pi)^{1/2}\frac{p}{1-p}\frac{\sigma}{\beta})]^{1/2} = \frac{t_0 - (1-p)\delta^* - M}{\sigma}. \tag{25}$$

In applying the inverse of $\phi(\cdot)$, it is crucial that the exogenous parameters lie in the admissible range: $0 < \frac{p}{1-p}\frac{\sigma}{\beta} < \frac{1}{(2\pi)^{1/2}}$. If $p = 0$, then self control is free and everyone can avoid discontentment with probability one by choosing $\delta = \infty$. If the parameters are outside the admissible range in the other direction, $\frac{1}{(2\pi)^{1/2}} < \frac{p}{1-p}\frac{\sigma}{\beta}$, then self control is too expensive relative to the β-weighted concern over the possibility of discontentment, and $\delta = 0$ is the expected utility maximizing choice. The admissibility condition for an interior δ^* can be re-expressed as:

$$p < \frac{1}{\frac{\sigma}{\beta}(2\pi)^{1/2} + 1}, \tag{26}$$

which implies, too, that $p < 1$ must hold. Otherwise, if $p \geq 1$, it would cost more than one unit of expected ordinary consumption to reduce the contentment threshold by one unit, which would be harmful on net to the decision maker: if μ falls by a magnitude greater than δ, then the probability of discontentment goes up, not down.

We can now use (25) to solve explicitly for δ^*:

$$\delta^* = \frac{1}{1-p}(t_0 - M + \sigma[-2\log((2\pi)^{1/2}\frac{p}{1-p}\frac{\sigma}{\beta})]^{1/2}), \tag{27}$$

provided all conditions on the exogenous parameters are met that guarantee a strictly interior maximizer.

Demand for self control as given by $\delta^*(\theta)$ turns out to be non-monotonic in p, and discontinuous in p on some subsets of the border regions of the admissible parameter set where the right-hand side of (27) is exactly zero or at its maximum $(M - \mu_0)/p$.

As illustrated in subsequent sections, small changes in price can cause erratic jumps in demand for self control, from zero to strictly positive levels, without continuously increasing as a function of p. The non-monotonic and sometimes highly unstable statics of self control with respect to exogenous variables θ are analyzed next and depicted with graphical analyses.

3 Consumer Response to Changes in the Decision Making Environment: Statics of $\delta^*(\theta)$

As new consumer finance regulations, programs and products are rolled out, a basic question is how consumers will respond to incentives used to encourage uptake of those new offerings. An incentive to take up such offerings would, in our model, amount to a reduction in p. Thus, the response of δ^* to p is perhaps the most important relationship to characterize. Before presenting graphs showing δ^* as a function of p, standard calculus-based expressions are stated to reveal both their insights and limitations. One issue with $\delta^*(\theta)$ as a function of p is discontinuity. In some regions of parameter space, $\delta^*(\theta)$ does not respond at all to changes in p. In other regions, however, very small changes in p cause optimal self control to discontinuously alternate between corner and interior quantities of self control, at which points $\frac{\partial \delta^*(\theta)}{\partial p}$ does not exist. The question of where in the parameter space corner solutions prevail and what happens along the boundary between interior and corner values of $\delta*$ are primary points of interest. The results reported in this section apply only to parameter values of θ at which $\delta^*(\theta)$ is interior.

According to formula (27), an exogenous increase in a person's (individual, community, culture, or country's) contentment threshold t_0, while holding material resources

M and all other exogenous variables constant, exerts a strictly positive effect on δ^*:

$$\frac{\partial \delta^*(\theta)}{\partial t_0} = \frac{1}{1-p} > 0. \tag{28}$$

When t_0 rises without any increase in resources available to finance consumption, then the decision maker faces increased risk of discontentment. The decision maker responds by re-allocating effort away from ordinary consumption and into self control.

One scenario that matches this prediction would be a comparison of religiosity in a number of poor countries between, say, 1980 and 2000. Increasing information (e.g., via internet) about consumption levels in rich countries whose national incomes were growing would exert upward pressure on t_0 even in countries not enjoying growth. The model implies that a rational response to such a scenario may not involve re-doubling efforts to catch up in terms of growth opportunities but rather re-allocate effort away from ordinary consumption and toward self control instead. Many religions provide, among many other services, technology for finding fulfillment without increasing consumption. If one thinks about the feudal economies under which organized Christianity developed, or the much earlier economies of Asia in which Buddhist philosophy flourished, the model would predict that these expansions of practices that teach moderation of desire would coincide with increasingly available signals available to poor people about the high levels of ordinary consumption enjoyed by others (interpreted as an upward shift of t_0).

The effect of greater wealth (e.g., advances in labor-saving technology that increase the value of one's endowment of effort) exerts an unambiguously negative effect on the demand for moderation of desire:

$$\frac{\partial \delta^*(\theta)}{\partial M} = -\frac{1}{1-p} < 0. \tag{29}$$

Therefore, in this model, self control is unambiguously an inferior good. This inferior good result hints at the negative income effects we will see below that are strong enough to produce upward-sloping demand curves.

Interestingly, if the contentment threshold and the real value of one's endowment of resources grow by equal amounts, the two effects cancel each other. Changes in contentment thresholds (i.e., aspirations or goals needed to achieve contentment) and changes in wealth only have effects on self control net of changes in the other. Shifts in the contentment threshold t_0 have an effect on self control when t_0 changes by more than M. And shifts in M are similarly important only if they do not track changes in t_0.

It is useful to denote the bracketed term in (27) as the function $h(y) \equiv [-2\log((2\pi)^{1/2}y)]^{1/2} = z$ (with y evaluated at $\frac{p}{1-p}\frac{\sigma}{\beta}$), mapping standard normal pdf values (y) into z on the positive real line. One notices that $h(y)$ is a positive valued and strictly decreasing function. Taking advantage of the inequality $h'(y) < 0$, it is straightforward to verify the intuitively obvious result that an increase in β, reflecting stronger subjective weight on the possibility of discontentedness, causes demand for self control to increase:

$$\frac{\partial \delta^*(\theta)}{\partial \beta} = -h'(\frac{p}{1-p}\frac{\sigma}{\beta})p(\frac{1}{1-p}\frac{\sigma}{\beta})^2 > 0. \tag{30}$$

The parameter σ represents the volatility or riskiness of ordinary consumption. Another interpretation of σ is as a proxy for wealth inequality orthogonal to effort, since the distribution of material standards of living will be more spread out in environments with large σ, even if everyone were to choose the same level of effort. To characterize the sign of $\frac{\partial \delta^*(\theta)}{\partial \sigma}$, it is useful to notice that $h'(y) = -\frac{1}{yh(y)}$ and apply this result to simplify the following expression:

$$\frac{\partial \delta^*(\theta)}{\partial \sigma} = \frac{1}{1-p}[h(\frac{p}{1-p}\frac{\sigma}{\beta})+\sigma h'(\frac{p}{1-p}\frac{\sigma}{\beta})\frac{p}{1-p}\frac{1}{\beta}] = \frac{1}{1-p}[h(\frac{p}{1-p}\frac{\sigma}{\beta})-[h(\frac{p}{1-p}\frac{\sigma}{\beta})]^{-1}]. \tag{31}$$

The sign of $\frac{\partial \delta^*(\theta)}{\partial \sigma}$ depends on whether $h(\frac{p}{1-p}\frac{\sigma}{\beta})$ is greater or less than 1. By definition, $h(\phi(z)) = z$ for positive z. Therefore, the sign of $\frac{\partial \delta^*(\theta)}{\partial \sigma}$ depends on whether $\frac{p}{1-p}\frac{\sigma}{\beta}$ is greater or less than $\phi(1) \approx 0.2420$. Keeping in mind that h is decreasing and therefore that $h(y) < 1$ if $y > \phi(1)$ and $h(y) > 1$ if $y < \phi(1)$, the direction of the desired effect can be characterized as:

$$\frac{\partial \delta^*(\theta)}{\partial \sigma} > 0 \text{ whenever } \frac{p}{1-p}\frac{\sigma}{\beta} < \phi(1), \text{ and } \frac{\partial \delta^*(\theta)}{\partial \sigma} < 0 \text{ whenever } \frac{p}{1-p}\frac{\sigma}{\beta} > \phi(1). \tag{32}$$

Thus, the effect of σ on $\delta^*(\theta)$ is in general non-monotonic. For moderate levels of σ and p, the effect is positive, indicating increased demand for self control when risk or inequality rises. The intuition for this is that, in low-risk low-inequality environments, a small increase in risk or inequality increases the probability of discontentment on a relatively steep range of the pdf ϕ, implying that increases in δ have proportionally large effects. In contrast, when beginning on a relatively flat range of the pdf farther than one standard deviation from the mean, adjustments to δ have a much weaker effect in terms of reducing the risk of discontentment.

Finally, the effect of p on $\delta^*(\theta)$ is, once again, non-monotonic, as the figures in the next section will show. We record the analytic expression here:

$$\frac{\partial \delta^*(\theta)}{\partial p} = \frac{1}{(1-p)^2}[t_0 - M + \sigma h(\frac{p}{1-p}\frac{\sigma}{\beta}) + \frac{1}{1-p}\frac{\sigma^2}{\beta}h'(\frac{p}{1-p}\frac{\sigma}{\beta})]. \tag{33}$$

This can be developed, once again, using $h'(y) = \frac{1}{yh(y)}$. It is unclear, however, whether this leads to any new insights about regions of parameter space in which demand for self control is upward- versus downward-sloping.

4 Demand for Self Control

Figure 1 presents nine demand curves for self control, plotting δ^* as a function of price for particular values of the other parameters. The x-axis is δ^*, the distance by which the individual chooses to shift the contentment threshold t downward from its default t_0. The y-axis is the full range of prices from zero to the upper bound given in

(26). The resource endowment M is normalized to 1 for all graphs. In Figure 1, t_0 is normalized to 1. In later figures, t_0 is varied up and down by two standard deviations. Figure 1 varies σ moving from top to bottom and β moving left to right. Moving down Figure 1 from top to bottom, the random component of ordinary consumption X is increasing as indicated by the values of $\sigma = 0.1$, 1.0 and 10.0. The x-axes are re-scaled in the three rows to better see the detail of the curves. Moving left to right, values of β are 0.5, 1, and 2, reflecting increasing subjective weight on the risk of discontentment.

Non-monotonicity of $\delta^*(\theta)$ as a function of p is evident in the upward-sloping (i.e., backward bending) regions of the upper center and upper right subfigures of Figure 1. The upper right demand curve in Figure 1 also shows extreme sensitivity of demand for self control to price near $p = 0.9$. In this range, quantity demanded can shift from 0 to the maximum amount possible in response to very small changes in price, suggestive of cultural or religious revolutions whose philosophies depend heavily on critiques of excessive focus on ordinary consumption and advocating the practice of strict self control. We leave it to future research to pursue the question of whether these large price effects in opposite directions along the same demand curve (as in the upper right subfigure of Figure 1) might provide an explanation for the emergence of religious and spiritual movements observed in the historical record. Their cautionary implication for programs and new product launches aiming to improve self control should be clear from the difficult-to-predict responses to price incentives in Figure 1. Unstable price effects on demand for self control are further elaborated upon below.

The response of consumption risk on self control, $\frac{\partial \delta^*(\theta)}{\partial \sigma}$, can also be seen to be non-monotonic in Figure 1, by reading off the quantity demanded corresponding to a price of 0.02 on the three subfigures along the left. In the topmost left subfigure corresponding to $\sigma = 0.1$ and $\beta = 0.5$, a price of 0.02 would correspond to a value of

δ^* of slightly more than 0.3. Moving down to the center left subfigure corresponding to $\sigma = 1$ and $\beta = 0.5$, a price of 0.02 would correspond to a value of δ^* more than 2, slightly less than 3. In the bottom left subfigure corresponding to $\sigma = 10$ and $\beta = 0.5$, however, a price of 0.02 would correspond to a value of δ^* of exactly 0. These three point evaluations for successively increasing values of σ show that, for the parameter values considered, $\delta^*(\theta)$ is an increasing function of σ when evaluated at $\sigma = 0.1$ but decreasing when evaluated at $\sigma = 1$. This non-monotonicity is not an artifact of scaling.

Moving from left to right in Figure 1, the weight on discontentment in the objective function is increasing, and demand for self control consequently shifts out to the right just as one would expect. The combination of low σ and large β produces demand curves with an upward-sloping portion mentioned above, reflecting large income effects on the demand for self control.

Figure 2 illustrates an analogous set of demand curves, this time with a default contentment threshold that is very easy to reach: $t_0 = M - 2\sigma$. This corresponds to an environment in which more than 95 percent of people will achieve contentment, even when allocating zero effort to self control (so that $t = t_0$). In this case, the possibility of discontentment is remote. Positive quantities of self control are nevertheless demanded, which reduces the chance of discontentment even more. No unusual features are seen when t_0 is already very easy to reach, with self control increasing monotonically as its price declines.

Figure 3 illustrates another analogous set of environments, this time with a default contentment threshold that is very difficult to reach: $t_0 = M + 2\sigma$. In this case, if zero self control were demanded, then risk of discontentment is more than 95 percent. Figure 3 shows that environments with low σ and high t_0, in which discontentment is virtually certain, produce demand curves with large upward-sloping regions. These strongly negative income effects are intuitive, because the high degree of certainty of discontentment on the convex portion of the pdf produces something tantamount to increasing returns to self control. The more self control is applied, the more productive self control becomes at reducing the probability of discontentment.

Figure 4 presents expenditure-price curves to visualize what fraction of total income is allocated to self control over the price range. Recalling that the budget constraint $\mu + p\delta = M$ with M normalized to 1, the quantity $100p\delta$ gives the percentage of the wealth endowment (net of exogenously given μ_0) allocated to self control. Figure 4 presents nine expenditure-price curves with the same parameter values as in Figure 1. These show that non-monotonic expenditures on self control—increasing and then decreasing expenditures on self control as prices increases from zero—is the rule rather than the exception. When price is very low, a large quantity of self control can be purchased for a small expenditure, moving the contentment threshold to the concave portion of the pdf of X. When price is very large, the sacrifice in terms of μ is so great that only a small portion of the endowment is allocated to self control. In the middle of the price range, however, a larger and larger share of wealth is allocated to self control reflecting movement along the convex (increasing-returns) portion of the pdf until the concave portion is reached, at which point total expenditures decline.

Finally, we document price discontinuities in the demand for self control that occur for some parameterizations. Discontinuity occurs when quantities demanded jump from zero to a substantial share of wealth, in response to a small movement in price. Such a discontinuity is illustrated in the three snapshots of the consumer's objective function depicted in the three subfigures of Figure 5. Figure 5 shows the univariate objective function $v(\delta)$ from equation (20) at three nearby values, $p = 0.70$, 0.72, and 0.74. The values of all other parameters are as in Figure 1, except for $t_0 = 0.5$ to reflect a scenario in which discontentment is defined by all levels of ordinary consumption more than half a standard deviation below the default mean of X, μ_0. This ensures that both competing motives—applying effort, on the one hand, to increase the mean of X and, on the other, to decrease the threshold at which discontentment is defined—are distinctly reflected in the objective function. The three curves representing $v(\delta)$ shift smoothly as price increases from 0.70 to 0.74. In the first and second subfigures of Figure 5, the maximizer δ^* is in both cases interior. However, moving from the second to third subfigure, the global maximizer of $v(\delta)$ jumps discontinuously, with an abrupt shift to the left-most corner, $\delta^* = 0$, seen in the third subfigure.

5 Summary and Discussion

When the cost of self control is sufficiently large, demand for self control is zero and the consumer chooses a path to happiness exclusively through the channel of ordinary consumption. When the price of self control falls below this critical threshold (which is a function of nearly all the exogenously given parameters in the environment), demand for self control responds systematically, although not always monotonically, to shifts in the parameters that describe the decision maker's environment. There are three globally monotonic effects on the demand for self control to report. First, self control is an inferior good, always decreasing in response to growing wealth, holding all other parameters, especially the decision maker's goals and aspirations, constant. More intense desire for ordinary consumption (i.e., raising the discontentment threshold t_0) increases demand for self control, as does the preference parameter β measuring the weight placed on the risk of falling short of one's contentment threshold. Self control is a monotonically decreasing function of price in highly uncertain environments but backward bending (with an upward sloping portion) when discontentment risk is large and nearly certain. The effect of uncertainty on demand for self control is, holding price of self control constant, non-monotonic, reflecting increasing and decreasing returns to self control that are a direct consequence of convex versus concave portions of the pdf of the discontentment risk function. This non-monotonicity would disappear given a triangle-shaped pdf without convex and concave segments in the domain.

Whether incentives to induce consumers to engage in more self control will be effective or not depends on where in the model's parameter space the decision making environment is. If the price of self control is way above the critical threshold where $\delta^*(\theta)$ cuts off to zero, then small inducements are predicted to have no effect. In other environments, $\frac{\partial \delta^*(\theta)}{\partial p}$ may be large or small, and of varying sign.

If the decision making environment happens to be on an upward sloping portion of the demand for self control curve, then regulation, programs and products offered with incentives that reduce the price of self control can have the unanticipated effect of decreasing demand for self control. For example, suppose we observe IRA retirement account limits raised by an act of the US Congress, which is interpreted as a reduction in p, and simultaneously observe Americans' savings rate to fall. In such an environment, high priced self control technologies induce greater uptake of self control than low priced offerings.[11]

The case most conducive for policy interventions to succeed at encouraging self control are the nearly flat but decreasing portions of the demand for self control curves. Figures 1, 3, 4 and 5 show a number of environments in which, starting from a parameter value θ at which $\delta^*(\theta) = 0$, a small reduction in p induces large shifts into self control. The discontinuous shift from zero to large expenditures on self control is seen moving from the right-most to the center subfigures of Figure 5.

[11] The widely followed Indian teacher of breathing exercises and philosophy, Sri Sri Ravi Shankar (who founded the Art of Living Foundation), says that it is important to charge a positive price for the services he offers, even though he is willing to give them away for free, because consumers value it more when its price is positive (Shankar, 2010).

The model features the possibility of rationally choosing to abandon self control, which challenges a basic definition of bounded rationality often put forward in the behavioral economics literature based on three bounds: on willpower, computational capacity, and self-interest. The model describes an outcome in which agents (and the cultures they populate) explicitly choose to abandon self control if the benefits of self control do not outweigh its costs. This abandonment of self control is a rational and predictable function of opportunities and costs in the environment. If one views different cultures' philosophical traditions through the lens of this model, one prediction (for appropriate values of θ) is that some cultures devote themselves wholly to ordinary consumption while others emphasize allocating effort to self control in a way that moderates desire for ordinary consumption. The mix of wealth, aspirations and a beneficial uncertainty that enables some to experience the good fortune of climbing upwards in the consumption distribution plays a crucial role creating the condition under which individuals or cultures might move from the corner solution of zero to positive levels of self control.

Although the model is behavioral in that it presumes willpower is costly and that preferences can be changed, the price theory analysis is neoclassical. We use a parsimonious and highly stylized neoclassical approach to the behavioral problem of choosing how much self control to acquire when making consumption decisions. The model is essentially the two-good, one-period consumer choice problem from undergraduate microeconomic textbooks, augmented with a term in the utility function that depends on the difference between the level of ordinary consumption chosen and a reference-point level of ordinary consumption needed to hit one's psychological target or, equivalently, avoid discontentment. Although it draws inspiration from the satisficing and goal setting literatures, the model is not a satisficing model because it wholly adopts the assumption of expected utility maximization and its implication of exhaustive search through the choice set rather than limited search as in the satisficing model.

The instability of demand for self control in even this simplest of models should perhaps give policy makers pause when predicting rates of uptake for new programs, conformity with new regulations, or forecasting demand for new financial services aimed at increasing self control. The concerns raised by this model do not touch on the difficult philosophical problems of paternalism or result from ambiguity about the meaning of "excessive" consumption. The model assumes that, after introspectively consulting our own experiences and temptations, it is possible for consumption to be excessive. Taking that process of reflection by the consumer on his or her own goals and means of avoiding discontentment, the model describes how consumers systematically use self control technology in their real world environments to reallocate resources away from ordinary consumption to other activities that raise utility by reducing risk of discontentment.[12]

The point we want to make is that technologies aimed at aiding consumer self control, because of the strong income effects they produce, can lead to counterintuitive effects of interventions intended to encourage demand for self control. Without conclusive evidence to the contrary, self control should perhaps be assumed to be non-monotonic in price and therefore highly unstable. The possibility of upward-sloping demand schedules for self control raises the specter of policy interventions with grossly unanticipated or otherwise disappointing effects.

This theoretical finding is consistent with empirical accounts of interventions aimed

[12] This rationalization of partial or total lack of self control draws inspiration from the work of Caplan (2001), whose model predicts that irrationality rationally accumulates where it is least costly while responding systematically to incentives in the environment. The rational irrationality framework is expanded by Beaulier and Caplan (2007). Issues raised here regarding normative behavioral economics were first raised by Berg (2003) and elaborated in the context of behavioral economics and paternalism in a model of social welfare maximization for a society of satisficers (Berg and Gigerenzer, 2007).

at improving self control that wound up resulting in shifts in the opposite direction as was anticipated. Pence (2001) documents how interventions to induce households to save more have elicited essentially zero response, consistent with a decision making environment in which comparative statics are uniformly zero. The paper by do Valle, Pieters, and Zeelenberg (2008) demonstrates counterintuitive self control effects that can be interpreted as resulting from an upward sloping portion of the demand for self control curve. Many have expressed more general skepticism about the wisdom of consumer protections founded on the premise that consumer behavior is pathologically biased. In some financial market models (e.g., Berg and Lien, 2005; Jerzmanowski and Nabar, 2008), biased beliefs or behavior leads to positive welfare effects, raising the possibility that loss of self control in some contexts can have positive externalities. In such environments, interventions aiming to de-bias individual behavior might reduce efficiency and, with it, social welfare.

The model introduced here also underscores that time inconsistency and the multiple-selves framework are not necessary for modeling self control. No time dimension is needed whatsoever. We view this static model of demand for self control as complementary to the insightful work on dynamical models of self control, especially Gifford's (2008) emotions-based approach and the work on adaptive aspirations by Güth (2010). We conclude that, for purposes of predicting consumer demand for products and services that aid self control, one would like to have a much richer empirical record, with numerous experiments in different locations and with different populations so that different values of the other parameters in the model might be controlled for. Without that data, one default position consistent with this model would prefer very simple regulatory rules that fall on corporations over those that impinge directly on the choice sets of consumers, which would likely produce more predictable results on the supply side by avoiding the unpredictability of non-monotonic demand for self control.

References

Ashraf, N., Karlan, D., and Yin, W. (2005). Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines. *Quarterly Journal of Economics*, 121, 635-672.

Ameriks, J., Caplin, A., Leahy, J., and Tyler, T. (2007). Measuring self-control problems. *American Economic Review, 97*(3), 966-972.

Bae, S., and Ott, A. (2008). The public economics of self control: Symposium in honor of Richard Abel Musgrave. *Journal of Economics and Finance, 32*(4), 356-367.

Barr, M., and Dokko, J. (2008). Paying to save: Tax withholding and asset allocation among low- and moderate-income taxpayers. U of Michigan Law & Economics, Olin Working Paper No. 07-22.

Battaglini, M., Benabou, R., and Tirole, J. (2005). Self control in peer groups. *Journal of Economic Theory,* 123(2), 105-134.

Beaulier, S. and Caplan, B. (2007). Behavioral economics and perverse effects of the welfare state. *Kyklos*, 60, 485-507.

Benartzi, S., & Thaler, R. (2007). Heuristics and biases in retirement savings behavior. *Journal of Economic Perspectives*, 21(3), 81-104.

Berg, N. (2003) .Normative behavioral economics, *Journal of Socio-Economics*, 32, 411-427.

Berg, N. and Gigerenzer, G. (2007). Psychology implies paternalism?: Bounded rationality may reduce the rationale to regulate risk-taking, *Social Choice and Welfare,* 28(2), 337-359.

Berg, N. and Gigerenzer, G. (2010). As-if behavioral economics: Neoclassical economics in disguise?, *History of Economic Ideas*, 18(1), 133-166.

Berg, N. and Lien, D. (2005). Does society benefit from investor overconfidence in the ability of financial market experts?, *Journal of Economic Behavior and Organization*, 58, 95-116.

Bertaut, C., Haliassos, M., and Reiter, M. (2009). Credit card debt puzzles and debt revolvers for self control. *Review of Finance*, 13(4), 657-692.

Bergstresser, D., and Poterba, J. (2004). Asset allocation and asset location: Household evidence from the Survey of Consumer Finances. *Journal of Public Economics*, 88(9-10), 1893-1915.

Bertaut, C., Haliassos, M., and Reiter, M. (2009). Credit card debt puzzles and debt revolvers for self control. *Review of Finance*, *13*(4), 657-692.

Beverly, S., and Sherraden, M. (1999). Institutional determinants of saving: Implications for low-income households and public policy. *Journal of Socio-Economics*, 28(4), 457-473.

Bruni, L., and Porta, P. L. (2007). *Handbook on the Economics of Happiness*. Cheltanham, UK; Northhampton, MA: Elgar.

Caplan, B. (2001). Rational ignorance versus rational irrationality. *Kyklos*, 54, 3-26.

Carrillo, J. (1998). Self-control, moderate consumption and craving. *CEPR Discussion Papers.*

Castronova, E., and Hagstrom, P. (2004). The demand for credit cards: Evidence from the Survey of Consumer Finances. *Economic Inquiry*, 42(2), 304-318.

DeJong, D., and Ripoll, M. (2006). Do self-control preferences help explain the puzzling behavior of asset prices?. Working Paper, University of Pittsburgh.

Dey, S., and Mumy, G. (2005). Determinants of borrowing limits on credit cards. Working Paper 05-7, Bank of Canada.

do Valle, R., Pieters, R., and Zeelenberg, M. (2008). Flying under the radar: Perverse package size effects on consumption self-regulation. *Journal of Consumer Research*, *35*(3), 380-390.

Drolet, A., Luce, M., and Simonson, I. (2009). When does choice reveal preference? Moderators of heuristic versus goal-based choice. *Journal of Consumer Research*, *36*(1), 137-147.

Duygan-Bump, B., and Grant, C. (2008). Household debt repayment behaviour: What role do institutions play?. *Economic Policy*, *24*(57), 107-140.

Emmons, W. (2005). Consumer-finance myths and other obstacles to financial literacy. Supervisory Policy Analysis Working Papers 2005-03, Federal Reserve Bank of St. Louis.

Ernst, F., Hogan, B., Vallas, M., Cook, M., and Fuller, D. (2009). Superior self-regulatory skills in African American college students: Evidence from alcohol and tobacco use. *Journal of Black Studies*, 40(2), 337-346.

Feldstein, M. (1995). College scholarship rules and private saving. *American Economic Review*, *85*(3), 552-566.

Fisher, P., and Montalto, C. (2010). Effect of saving motives and horizon on saving behaviors. *Journal of Economic Psychology*, *31*(1), 92-105.

Fryer, R. (2010). Financial incentives and student achievement: Evidence from randomized trials. NBER Working Papers 15898.

Fudenberg, D., and Levine, D. (2006). A dual-self model of impulse control. *American Economic Review*, *96*(5), 1449-1476.

Gamber, W. (2003). Antebellum reform: Salvation, self-control, and social transformation. *Charity, Philanthropy, and Civility in American History* , pp.129-153, Cambridge; New York and Melbourne.

George, D. (2001). *Preference Pollution: How Markets Create The Desires We Dislike*. Ann Arbor: University of Michigan Press.

Gifford, A. (2002). Emotion and self-control. *Journal of Economic Behavior and Organization*, 49, 113-130.

Gigerenzer, G., and Selten, R. (2002). *Bounded Rationality: The Adaptive Toolbox*. Cambridge and London: MIT Press.

Güth, Werner, (2010). Satisficing and (un)bounded rationality--A formal definition and its experimental validity. *Journal of Economic Behavior and Organization*, *73*(3), 308-316.

Hastings, J., and Washington, E. (2010). The first of the month effect: Consumer behavior and store responses. *American Economic Journal: Economic Policy*, *2*(2), 142-162.

Hilton, D. (2001). The psychology of financial decision-making: Applications to trading, dealing, and investment analysis. *Journal of Psychology and Financial Markets*, *2*(1), 37-53.

Jain, S. (2009). Self-control and optimal goals: A theoretical analysis. *Marketing Science*, *28*(6), 1027-1045.

Jerzmanowski, M., and Nabar, M. (2008). The welfare consequences of irrational exuberance: Stock market booms, research investment, and productivity. *Journal of Macroeconomics*, *30*(1), 111-133.

Kan, K. (2007). Cigarette smoking and self-control. *Journal of Health Economics*, *26*(1), 61-81.

Karlsson, N. (2003). Consumer self-control strategies: An empirical study of their structure and determinants. *Journal of Consumer Policy*, 26(1), 23-41.

Kawaja, M. (1969). The economic effects of regulation: A case study of the consumer finance industry. *Southern Economic Journal*, 35(3), 231-238.

Kerr, S., and Dunn, L. (2008). Consumer search behavior in the changing credit card market. *Journal of Business and Economic Statistics*, *26*(3), 345-353.

Kim, J. (2006). Hyperbolic discounting and the repeated self-control problem. *Journal of Economic Psychology, 27*(3), 344-359.

Kumru, C., and Thanopoulos, A. (2008). Social security and self control preferences. *Journal of Economic Dynamics and Control*, *32*(3), 757-778.

Lai, C. (2006). Determinants of portfolio efficiency losses in US self-directed pension accounts. *Journal of Family and Economic Issues*, *27*(4).

Laibson, D., Repetto, A., and Tobacman, J. (1998). Self-control and saving for retirement. *Brookings Papers on Economic Activity*, 1, 91-196.

Laran, J. (2010). Choosing your future: Temporal distance and the balance between self-control and indulgence. *Journal of Consumer Research*, *36*(6), 1002-1015.

Lee, L., and Ariely, D. (2006). Shopping goals, goal concreteness, and conditional promotions. *Journal of Consumer Research*, *33*(1), 60-70.

Lin, B., and Yen, S. (2008). Consumer knowledge, food label use and grain consumption in the US. *Applied Economics*, *40*(4-6), 437-448.

Lusardi, A., and Mitchell, O. (2006). Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth. Working Paper, Wharton School.

Oaten, M., and Cheng, K. (2007). Improvements in self-control from financial monitoring. *Journal of Economic Psychology*, *28*(4), 487-501.

O'Donoghue, T., and Rabin, M. (2002). Addiction and Present-Biased Preferences. Working Paper, Cornell University; Economics Department; University of California, Berkeley.

Otto, P., Davies, G., and Chater, N. (2007). Note on ways of saving: Mental mechanisms as tools for self-control?. *Global Business and Economics Review*, *9*(2-3), 227-254.

Pence, K. (2001). Essays on Government Policy and Household Financial Decisions. Thesis, University of Wisconsin.

Ramanathan, S., and Williams, P. (2007). Immediate and delayed emotional consequences of indulgence: The moderating influence of personality type on mixed emotions. *Journal of Consumer Research*, *34*(2), 212-223.

Sally, D. (2000). I, too, sail past: Odysseus and the logic of self-control. *Kyklos*, *53*(2), 173-200.

Samwick, A. (2006). Saving for retirement: Understanding the importance of heterogeneity. *Business Economics*, *41*(1), 21-27.

Semboja, H. (2004). Small is beautiful, but growth is inevitable: Experiences of apex institutions in Senegal and Tanzania. *Development Southern Africa*, *21*(5), 867-878.

Servon, L., and Kaestner, R. (2008). Consumer financial literacy and the Impact of online banking on the financial behavior of lower-income bank customers. *Journal of Consumer Affairs*, *42*(2), 271-305.

Shankar, S. S. R. (2010). *The Interview*. BBC, August 14, 2010. http://www.bbc.co.uk/iplayer/episode/p0090py7/The_Interview_14_08_2010_Sri_Sri_Ravi_Shankar/

Simon, H. A. (1956). Rational choice and the structure of environments. *Psychological Review* 63, 129-138.

Sprenger, C., and Stavins, J. (2008). Credit card debt and payment use. Working Paper 08-2, Federal Reserve Bank of Boston.

Smith, T. (2002). The McDonald's equilibrium: Advertising, empty calories, and the endogenous determination of dietary preferences. *Social Choice and Welfare*, 23, 383-413.

Starr, M. (2007). Saving, spending, and self-control: Cognition versus consumer culture. *Review of Radical Political Economics*, *39*(2), 214-229.

Sugden, R. (2004). The opportunity criterion: Consumer sovereignty without the assumption of coherent preferences. *American Economic Review*, 94(4), 1014-1033.

Thaler, R., and Benartzi, S. (2004). Save more tomorrow: Using behavioral economics to increase employee saving. *Journal of Political Economy*, *112*(1), S164-S187.

Thaler, R. H., and Shefrin, H. M. (1981). An economic theory of self-control. *Journal of Political Economy*, 89, 392-406.

Trudel, R., and Cotte, J. (2008). The effect of regulatory focus on satisfaction. *Advances in Consumer Research - European Conference Proceedings*, *8336*. Retrieved from Business Source Complete database.

Wertenbroch, K. (1998). Consumption self-control by rationing purchase quantities of virtue and vice. *Marketing Science*, *17*(4), 317-337.

White, M. (2009). Bankruptcy: Past puzzles, recent reforms, and the mortgage crisis. *American Law and Economics Review*, *11*(1), 1-23.

Vohs, K., and Faber, R. (2007). Spent resources: Self-regulatory resource availability affects impulse buying. *Journal of Consumer Research*, *33*(4), 537-547.

Yao, R., Hanna, S., and Lindamood, S. (2004). Changes in financial risk tolerance, 1983-2001. *Financial Services Review*, *13*(4), 249-266.

Zhang, Y., Huang, S., and Broniarczyk, S. (2010). Counteractive construal in consumer goal pursuit. *Journal of Consumer Research*, *37*(1), 129-142.

Figure 1: Demand for Self Control

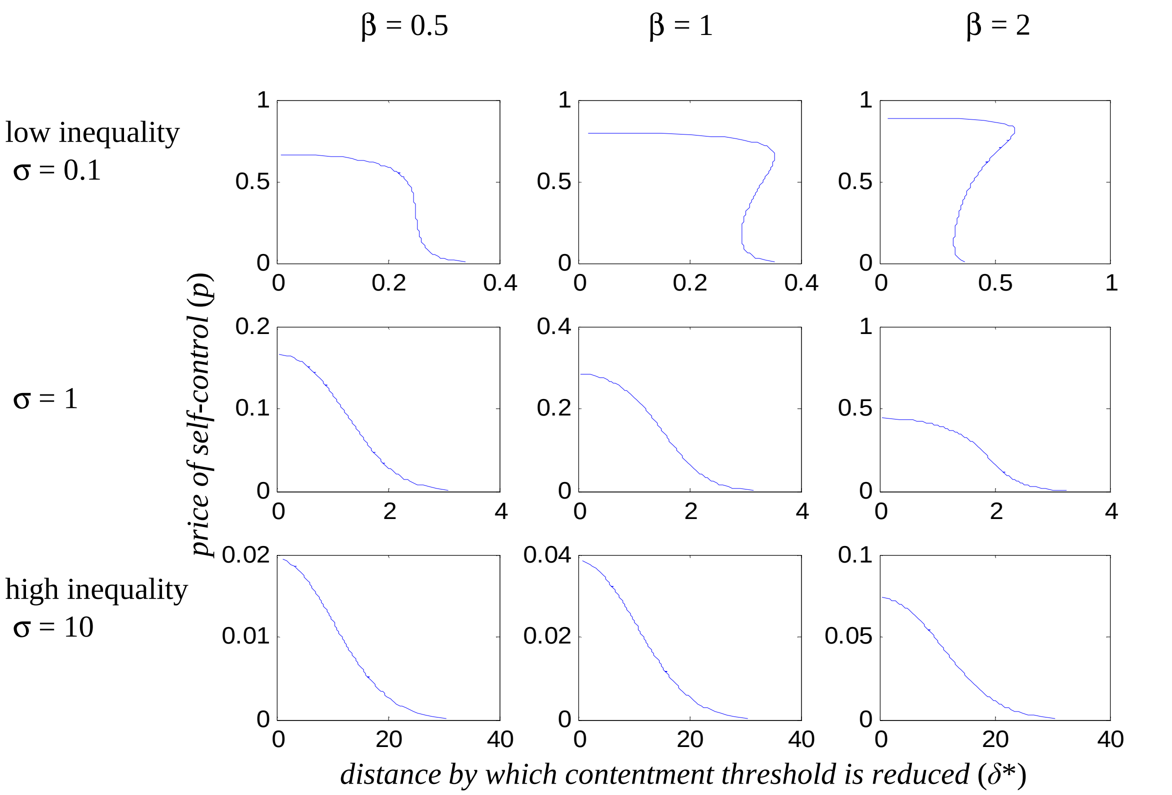

The demand curves above plot expected utility maximizing values of δ over prices ranging from zero to the upper bound in (26), with t_0 = 1 and M = 1.

Figure 2: Demand for Self Control with Easy-to-Reach Contentment Threshold Default ($t\_0 = M - 2\sigma$)

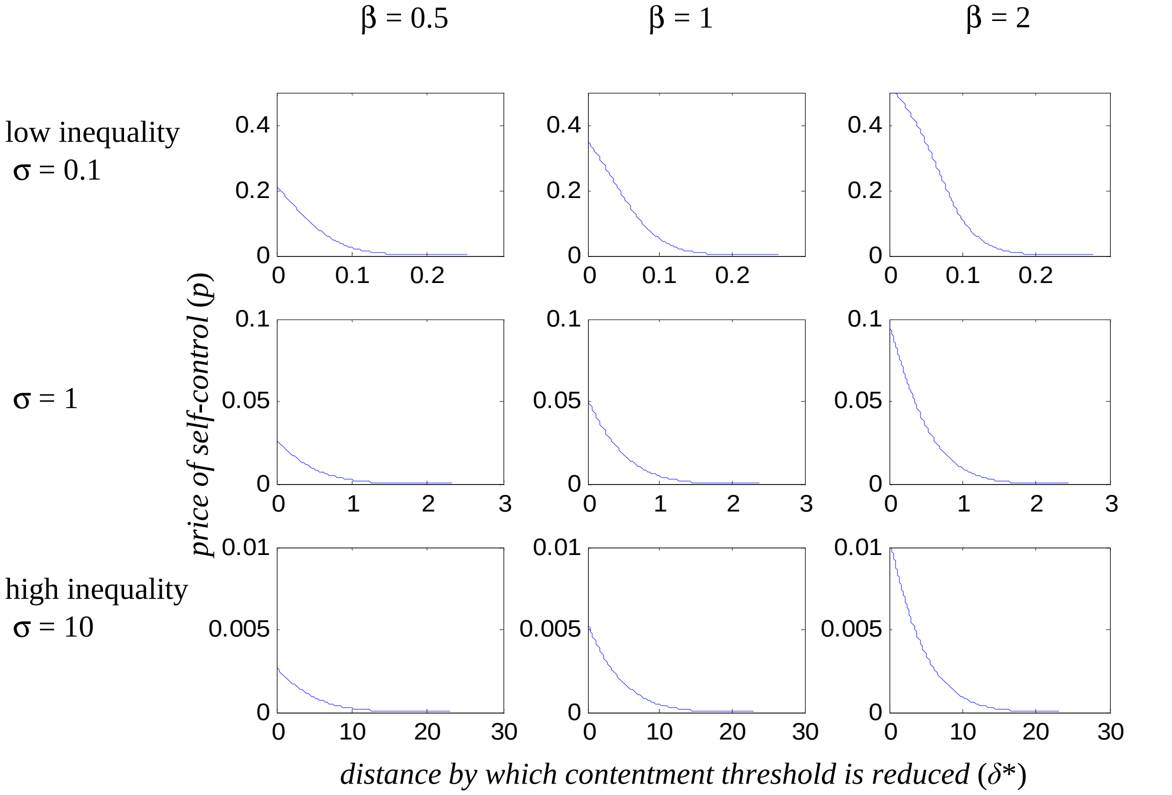

Figure 3: Ordinary Demand for Control with Hard-to-Reach Contentment Threshold Default ($t\_0 = M + 2\sigma$)

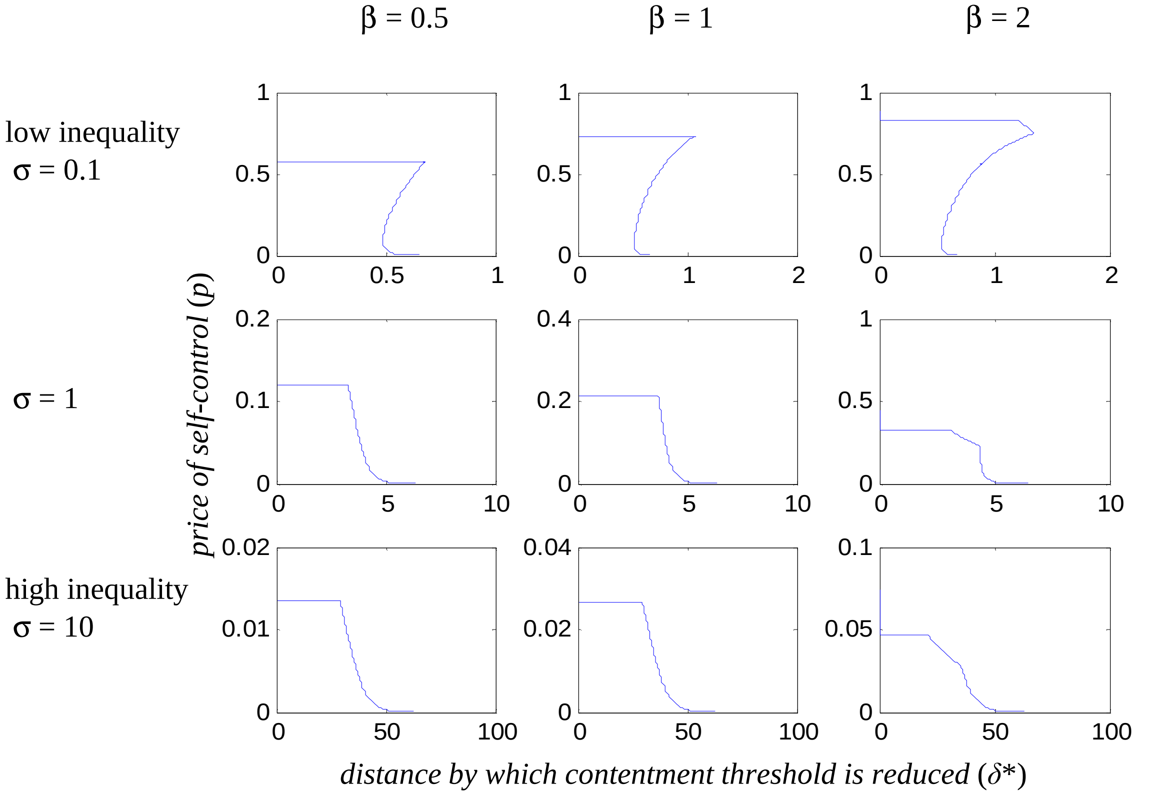

Figure 4: Percentage of *M* Expended on Self Control as a Function of Price

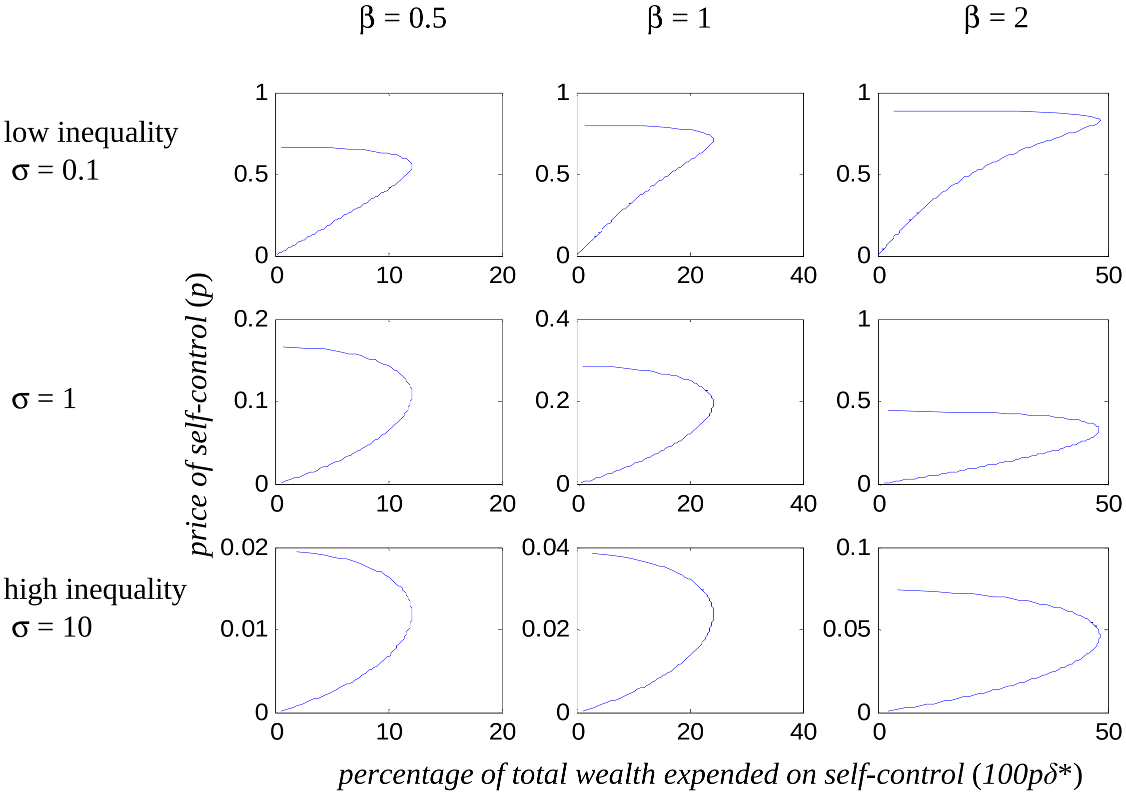

Figure 5: Discontinuity of δ*: Large Change in Global Maximizer Following Small Change in Price